

A curious dialogical logic and its composition problem

Sara L. Uckelman

(sara.uckelman@asia-europe.uni-heidelberg.de)*

Exzellenzcluster “Asia and Europe”, Karl Jaspers Centre for Advanced Transcultural Studies, Ruprecht-Karls Universität Heidelberg

Jesse Alama (j.alama@fct.unl.pt)

Center for Artificial Intelligence, New University of Lisbon

Aleks Knoks (aknoks@umd.edu)

Department of Philosophy, University of Maryland

Abstract. Dialogue semantics for logic are two-player logic games between a Proponent who puts forward a logical formula φ as valid or true and an Opponent who disputes this. An advantage of the dialogical approach is that it is a uniform framework from which different logics can be obtained through only small variations of the basic rules. We introduce the *composition problem* for dialogue games as the problem of resolving, for a set S of rules for dialogue games, whether the set of S -dialogically valid formulas is closed under *modus ponens*. Solving the composition problem is fundamental for the dialogical approach to logic; despite its simplicity, it often requires an indirect solution with the help of significant logical machinery such as cut-elimination. Direct solutions to the composition problem can, however, sometimes be had. As an example, we give a set N of dialogue rules which is well-justified from the dialogical point of view, but whose set N of dialogically valid formulas is both non-trivial and non-standard. We prove that the composition problem for N can be solved directly, and introduce a tableaux system for N .

Keywords: composition problem, dialogical logic, tableaux

* The first two authors were partially funded by the FCT/NWO/DFG project “Dialogical Foundations of Semantics” (DiFoS) in the ESF EuroCoRes programme LogICCC (FCT LogICCC/0001/2007; LogICCC-FP004; DN 231-80-002; CN 2008/08314/GW).

Keywords: composition problem, dialogical logic, tableaux

1. Introduction

Dialogical logic was introduced by Lorenzen in the 1950s and developed by Lorenzen and Lorenz in the 1970s [23, 24]. The basis of dialogical logic is a two-player game (dialogue) between a Proponent (**P**) who attempts to show via a dialogue-game that a formula φ is valid; the other player, Opponent (**O**), disputes this. As with other logic games [14], less attention is paid to actual plays of dialogue games than to the tree of all possible ways the game could go, given an initial formula φ ; of particular interest is the existence of a winning strategy for **P**, which specifies how **P** can reply to any move of **O** in such a way that **P** can win. Lorenz claimed that Lorenzen’s dialogue games offer a new type of semantics for intuitionistic logic **IL** and asserted the equivalence between dialogical validity (defined in terms of winning strategies for **P**) and intuitionistic derivability [20, 21]. Lorenz’s proof contained some gaps, and later authors sought to fill these gaps; the first complete appeared in [9].¹

Dialogue games are not restricted to intuitionistic logic. By modifying the rules, these games can also provide semantics for both classical logic **CL** [1, 10] and other non-classical logics such as paraconsistent, connexive, modal and linear logics [16, 34]. Semantics for these logics, obtained by extending and adapting the original formulation of dialogue games, are achieved either by adding rules for new connectives, which are called particle rules, or by changing or adding rules governing the dialogues as wholes, called structural rules. Of course, one cannot change the rules randomly and expect to retain something sensible. It is a question of philosophical importance what types of changes to the rules can be well motivated and justified, and it is important to note that what is at issue is the justification of a ruleset **S**, and not a justification of the resulting set of **S**-valid formulas. Except at the risk of being *ad hoc*, one cannot adequately justify a set of dialogical rules solely on the basis that the ruleset characterizes some well-known logic, even though, as Krabbe notes, such a characterization is always welcome: “At the very least, [students of dialogue logic] want [their systems] to yield some well-known logic, so that they can prove completeness and correctness and get their papers published” [18, p. 35]. The point is that we cannot

always antecedently know that what we will obtain from a particular combination of rules is going to be “acceptable” (under some criteria of acceptability), and thus we must locate the problem of justification at the input (i.e., the ruleset) rather than the output (i.e., the set of dialogically valid formulas). Instead, there must be appeal to additional means of justification. Again, as Krabbe notes, “[a] problem for some motivations [in discussions of acceptable structural rules] was that it was felt to be highly desirable to end up with a system that yielded a respectable logic... but that, at least by some, it was not seen as permissible to let this desire be part of the motivation” [19, p. 694].

However, while many authors have recognized this, to date there has been little philosophical discussion of the problem of criteria for what constitutes a “good” or “justified” set of rules, and hence on what types of changes to a set of rules can and cannot be allowed.² The fact that there is no clear principled restriction on how dialogical rules can be modified naturally raises the question of when the set of S -valid formulas, for a particular set S of dialogical rules, actually corresponds to a logic. To answer this question, we need to identify properties that a set of formulas must have in order for it to be called a logic. One such desirable property is that the set be closed under *modus ponens*: If φ and $\varphi \rightarrow \psi$ are S -dialogically valid, then so should ψ be; we discuss others in §3. In this paper, we consider a dialogical ruleset which arises via a rather natural modification of a well-known ruleset, but our method of modification differs from previous attempts: Instead of changing a rule or adding to the set of rules, we *remove* a specific structural rule. We then investigate the properties of the resulting set of formulas which are dialogically valid according to this ruleset.

The plan of the paper is as follows. In the next section, we provide an introduction to (propositional) dialogical logic and give a dialogical definition of a new sub-classical propositional logic which we call N . In §3 we introduce the *composition problem* for a set of dialogical rules, relating it to the problem of showing that the set of formulas is a logic. In §4, we prove a number of results which lead to a positive solution to the composition problem for N . Then, in §5, we prove properties about N , locating it within the universe of known propositional logics. §§6 and 7 are devoted to a tableau-based characterization of N . We conclude in §8.

2. Dialogical logic

Developments in dialogical logic after Lorenzen and Lorenz’s seminal works can be roughly divided into two different trends: the Lille-school

of dialogues in the tradition of Rahman and Rückert on the one hand, and the proof-theoretic tradition of Felscher on the other. These traditions differ not in the underlying framework within which they work but in how the dialogues are presented. In Lille-style dialogues, the emphasis in the presentation is on “rounds”, that is, pairs of moves consisting of an attack and the corresponding defense, if possible, with the attack opening the round and the defense closing it. In this approach, the focus is on individual dialogues rather than winning strategies. Moreover, in the Lille tradition dialogues are generally restricted to be finite with the help of numerical bounds on the number of times the two players may repeat themselves. By contrast, in the Felscher tradition such bounds are absent, with the result that dialogues may be infinite. For more information on the Lille tradition, see [16, 31].

In the tradition of Felscher (who gave the first rigorous proof of correspondence between a particular set of dialogical validity and validity in intuitionistic logic [9]), dialogues are presented in such a way that makes clear their essential similarity to semantic tableaux or (Gentzen-style) sequent calculi. Such an approach simplifies completeness proofs for various dialogical logics due to the close connection between the dialogues and the proof theory. Dialogues building on Felscher’s notation (or showing strong similarity to his) can be found in [11, 12, 17, 36, 37, 40], among others.

In this paper we follow the tradition of Felscher, using [9] as our basis because it allows more perspicuous and direct conversions between dialogue strategies, tableaux, and sequent proofs and, thus, lends itself better to proving rigorous results in a straightforward fashion. Consequently, our results are not immediately applicable without modification to Lille-style dialogues. However, given that the two schools are rooted in the same approach to logic via dialogues, it should be possible for our results about Felscher-style dialogues to be adapted to Lille-style dialogues.

We work with a propositional language; formulas are built from atoms and \neg , \vee , \wedge , and \rightarrow . As is customary, lowercase Roman letters are used as atomic variables, and Greek letters are used as formula variables. In addition to formulas, there are the three so-called *symbolic attack* expressions, $?$, \wedge_L , and \wedge_R , which are distinct from all the formulas and connectives. Together formulas and symbolic attacks are called *statements*; they are what is asserted in a dialogue game.

The rules governing dialogues are divided into two types. *Particle rules* (also known as argumentation forms) define the meanings of connectives in a local fashion and say how formulas can be attacked and defended depending on their main connective. *Structural rules* operate globally and define what sequences of attacks and defenses count as

Table I. Particle rules for dialogue games

Assertion	Attack	Response
$\varphi \wedge \psi$	\wedge_L	φ
	\wedge_R	ψ
$\varphi \vee \psi$?	φ or ψ
$\varphi \rightarrow \psi$	φ	ψ
$\neg\varphi$	φ	—

dialogues, thus giving a global meaning to the connectives. Different logics can be obtained by modifying either set of rules.

The standard particle rules are given in Table I. According to the first row, there are two possible attacks against a conjunction: The attacker specifies whether the left or the right conjunct is to be defended, and the defender then continues the game by asserting the specified conjunct. The second row says that there is one attack against a disjunction; the defender then chooses which disjunct to assert. The interpretation of the third row is straightforward. The fourth row says that there is no way to defend against the attack against a negation; the only appropriate “defense” against an attack on a negation $\neg\varphi$ is to continue the game with the new information φ .

These notions can be made precise as follows.

Definition 1. A *signed expression* is a pair $\langle X, e \rangle$ where e is a statement and X is either **P** or **O**. A signed expression is said to be **P**-signed if its first component is **P** and **O**-signed if its first component is **O**. We sometimes call a signed expression $\langle X, e \rangle$ an *X-position*.

Let δ be a sequence (that is, a function from an initial segment of ω) of signed expressions and let η be a function for which:

- $\text{dom}(\eta) = \text{dom}(\delta) \setminus \{0\}$, and
- for every n in $\text{dom}(\delta)$, the value $\eta(n)$ is a pair $[m, Z]$, where m is a natural number less than n and Z is either “A” (attack) or “D” (defend).

Given such functions δ and η , the pair (δ, η) is a *dialogue* if it satisfies the three conditions:

1. If n is even, then $\delta(n)$ is a **P**-signed expression and if $\delta(n)$ is odd, then $\delta(n)$ is an **O**-signed expression.

2. If $\eta(n) = [m, A]$, then $\delta(m)$ is a non-atomic formula and $\delta(n)$ is an attack upon $\delta(m)$ according to the particle rules.
3. If $\eta(n) = [m, D]$, then there exists a natural number $k < n$ such that $\eta(m) = [k, A]$, and $\delta(n)$ is a defense against the attack $\delta(m)$ according to the particle rules.

If $\delta(0)$ is $\langle P, \varphi \rangle$, we say that the dialogue (δ, η) *commences with* φ .

These minimal conditions say only that play alternates between **P** and **O** (starting with **P** at move 0), and that every move (except the initial assertion $\delta(0)$) is either an attack or a defense against some earlier assertion.

Definition 2. A move in a dialogue is *assertive* if its assertion is a formula, that is, not a symbolic attack. An assertive move is *unattacked* if it is never attacked.

Further constraints on the development of a dialogue are given by the structural rules (see Def. 7 for a well-known set of structural rules). Because we are interested in working with different sets of structural rules, in the definitions below we abstract away from particular rulesets. These several definitions, and some of the discussion in §3, are thus parameterized by a set S of structural rules. Later we shall be interested in a handful of concrete rulesets.

Definition 3. Given a set S of structural rules, an *S-dialogue* for a formula φ is a dialogue commencing with φ that adheres to the rules of S . **P** *wins* an *S-dialogue* (δ, η) if there is a $k \in \mathbb{N}$ such that $\text{dom}(\delta) = [0, 2k]$ and there is no proper extension of (δ, η) , that is, there is no signed expression $\langle X, e \rangle$ and no natural number n such that δ could be extended to the domain $[0, 2k + 1]$ with the new value $\langle X, e \rangle$, with η likewise extended to have the value $[Z, n]$ at $2k + 1$.

According to this definition, if it is possible for the dialogue to continue (i.e., it is possible for one player to move), then neither player has yet won; the game proceeds as long as moves are available. In some settings, it may be preferable to enforce that any individual dialogue is finite, but it is not necessary, and we do not require that dialogues are finite here.

Definition 4. A *branch* of a rooted tree is a maximal totally ordered set of nodes that includes the root, where order is understood as the immediate ancestor relation. The *S-dialogue tree* $T_{S, \varphi}$ for a formula φ is the rooted tree satisfying the conditions:

- Every branch of $T_{S,\varphi}$ is an S -dialogue for φ ;
- Every S -dialogue for φ occurs as a branch of $T_{S,\varphi}$.

Remark. S -dialogue trees for non-atomic formulas can be quite complex, and indeed it often happens that some branches are infinite. If all branches are infinite, neither player wins.

Definition 5. An S -winning strategy s for \mathbf{P} for φ is a rooted subtree of $T_{S,\varphi}$ such that:

1. The root of s is the root of $T_{S,\varphi}$;
2. Every branch of s is an S -dialogue won by \mathbf{P} ;
3. If k is odd and a is a depth- k node of s , then a has exactly one child;
4. If k is even and a is a depth- k node of s , then a has the same children as does the image of a in $T_{S,\varphi}$.

Remark. Instead of saying “winning strategy for \mathbf{P} ” we simply say “winning strategy”.

This definition says that a winning strategy for \mathbf{P} is, in effect, a kind of function saying how \mathbf{P} should move given any move by \mathbf{O} and guarantee that he will win. Condition (1) requires that the strategy begins at the beginning. By Condition (2), the nodes of a winning strategy are all moves in a dialogue game and that all ways of playing according to the strategy end with a win for \mathbf{P} . Conditions (3) and (4) say that \mathbf{P} needs to have a unique response to any move \mathbf{O} could make in any of the dialogues that occur as branches in the strategy.

Dialogue games can be used to capture notions of validity.

Definition 6. For a set S of dialogue rules and a formula φ , the relation $\models_S \varphi$ means that \mathbf{P} has an S -winning strategy for φ . If $\not\models_S \varphi$, then we say that φ is S -invalid.

Note that, like proof-theoretic characterizations of validity, dialogue validity is an existential notion, unlike model-theoretic notions of validity, which are universal.

We now consider some example rulesets.

Definition 7. The ruleset D is comprised of the following structural rules [9, p. 220]:

- (D10) **P** may assert an atomic formula only after it has been asserted by **O** before: If $\delta(n) = \mathbf{P}a$ and a is atomic, then there exists $m < n$ such that $\delta(m) = \mathbf{O}a$.
- (D11) If p is an X -position, and if at $p - 1$ there are several open attacks made by Y , then only the *latest* of them may be answered at p : If $n(p) = [n, D]$ and if $n < j < p$, $j - n = 0$, $\eta(j) = [i, A]$, then there exists q such that $j < q < p$, $\eta(q) = [j, D]$.
- (D12) An attack may be answered at most once: For every n there exists at most one p such that $\eta(p) = [n, D]$.
- (D13) A **P**-assertion may be attacked at most once: If m is even, then there exists at most one n such that $\eta(n) = [m, A]$.

The main result of Felscher is the following:

THEOREM 1 (Felscher). *A formula φ is intuitionistically valid iff $\vDash_{\mathbf{D}} \varphi$.*

The proof goes by converting deductions in an intuitionistic sequent calculus to **D**-winning strategies (via tableaux), and vice versa. (The conversions are computable.)

Definition 8. The ruleset **D + E** is **D** plus the following rule:

- (**E**) **O** can react only upon the immediately preceding **P**-move: If n in $\text{def}(\delta)$ is odd, then $\eta(n) = [n - 1, Z]$, $Z = A$ or $Z = D$.

As Felscher notes, **E** implies **D13**, and, for odd p or n , also **D11** and **D12**. What is surprising is that we have that $\vDash_{\mathbf{D}} \varphi$ iff $\vDash_{\mathbf{D+E}} \varphi$ [9, p. 221]; we can change the set of moves available to **O** without changing the set of dialogically valid formulas.³

A subset of **D + E**, dropping rules **D11** and **D12** (though, again, the presence of **E** ensures that the effect of **D11** and **D12** partly remains), generates classical logic [1].

Definition 9. The ruleset **CL** is **D10 + D13 + E**.

One motivation which gave rise to results in this paper was the role of **E** in **CL** (the precise role of this rule is discussed in more detail in [40], and is not covered here). Adding **E** to the ruleset **D** did not change the set of validities: Both **D** and **D + E** correspond to intuitionistic logic. A natural question then is whether **E** is redundant in the same way in **CL** as it is in **IL**, or, more specifically, whether **D10 + D13 = D10 + D13 + E**.⁴ A major result of our paper is to show that this identity does not hold. Our counterexample is the logic **N**.

Definition 10. Let the ruleset $N = D10 + D13$. The logic N is the set of formulas for which \mathbf{P} has a winning N -strategy.

Surprisingly, even though it is generated by a natural, well-motivated, and straightforward modification of a ruleset, N will turn out to be radically different from classical logic, and we show that it diverges considerably from intuitionistic logic as well. As a result, it calls into question the utility of dialogue games as logical foundations without clear, principled, and non-*ad hoc* guidelines for modifying rulesets. Before we see this, however, we must first address what it means for a set of formulas to be a logic.

3. The composition problem

In this section we discuss properties that one might wish a set of formulas to have in order to be called a logic. The most important property, in our opinion, is closure under *modus ponens*, and we define the problem of determining whether a set of dialogical rules gives rise to an appropriately closed set:

Definition 11. (Composition problem) The *composition problem* for a set of dialogue rules S asks whether the set S of formulas φ for which \mathbf{P} has a winning strategy in the S -dialogue game for φ is closed under *modus ponens*.

That is, if \mathbf{P} has a winning S -strategy for φ and one for $\varphi \rightarrow \psi$, can we prove that \mathbf{P} has one for ψ ? A related problem is that of strategy composition:

Definition 12. (Strategy composition problem) The *strategy composition problem* for a set of dialogue rules S asks, given winning S -strategies for \mathbf{P} for formulas φ and $\varphi \rightarrow \psi$, whether we *compose* these strategies into one for ψ .

Clearly, a positive answer to the strategy composition problem will also be a positive answer to the more general problem, but the reverse is not the case: It may be possible that some set S of formulas is closed under *modus ponens*, but the winning strategies which generate the set are not composable. That is, a positive answer to the composition problem combined with a negative answer to the strategy composition problem would indicate the non-constructivity of the positive answer.

The composition problem is closely linked to what it conventionally means for a set of sentences to be a logic. Traditionally, a ‘logic’ has

been defined as a set of formulas closed under *modus ponens* and unrestricted uniform substitution [38]. However, even though the “regular closure under uniform substitution is an ordinary feature of many logics” [25, p. 221], it is not a universally-accepted feature, and requiring it is too strong, as it would force us to exclude from the category of ‘logic’ many systems that are generally agreed to be (non-classical) logics, for example: public announcement logic (PAL) [15]⁵, certain types of connexive logics (see, e.g., [27]), the inconsistency-adaptive logic LPM [28, p. 225, fn. 10], certain paraconsistent logics [32], approximated logics [13, p. 210], and indeed, nonmonotonic logics in general.⁶ Even if one were to insist on the presence of unrestricted uniform substitution, there are logics whose dialogical characterizations do not validate unrestricted uniform substitution, such as connexive logic [26], [34, §4.2] and relevance logic [34, §3.3].

Furthermore, as Marcos notes, “It has in fact often been assumed in the literature that *closure under uniform substitution* of atoms by complex sentences is presumed by the received (Aristotelian?) notion of ‘logical form’. But this needs not be so. For a start... no non-trivial extension of the consequence relation associated to classical logic (such as a supra-classical non-monotonic logic) can be defined unless uniform substitution is abandoned” [25, pp. 221]. This is related to the fact that the definition of logical consequence for a particular logic or theory can be either substitutional or not; that is, either the valid consequences are schemata or not.⁷

However, to avoid triviality, *some* constraint that goes beyond mere closure under *modus ponens* is required for an arbitrary set of atomic formulas to be called a logic.⁸ One natural constraint is to require the existence of at least one type of non-trivial restricted uniform substitution. We return to this issue in §5 below.

Thus, we deviate from the Tarskian definition by weakening the second requirement:

Definition 13. Given a language \mathcal{L} , a *logic* is a set L of \mathcal{L} -formulas which is closed under *modus ponens* (that is, if $\varphi \in L$ and $\varphi \rightarrow \psi \in L$, then $\psi \in L$ as well) and at least one form of uniform substitution (that is, there is some non-trivial function σ from atomic letters to formulas such that if $\varphi \in L$, then $\sigma(\varphi) \in L$).

A uniform solution to the composition problem for a wide range of rulesets S seems unrealistic. In specific cases, a single counterexample of course suffices for a negative solution, but proofs of positive solutions generally require heavy proof-theoretic machinery. Given the above-mentioned correspondences (completeness theorems) between dialogical validity and validity in various logics, there are a number of

rulesets S for which a positive solution to the composition problem for S holds, since the sets of formulas valid in intuitionistic logic, classical logic, connexive logic, various modal logics, etc., are all closed under *modus ponens*. However, the relevant completeness proofs in general use significant amount of non-dialogical machinery, specifically translations of dialogical strategies into derivations in some appropriate cut-free proof theory. One often-utilized method for proving a positive solution to what we are calling the composition problem is to give a correspondence between winning strategies and proofs or tableaux in proof-theory known to admit cut elimination [9, 12]. Thus, these correspondences require that one already have a proof theory for the target logic in question and that this proof system admits cut elimination; in many cases one or both of these may be lacking. As a result, in this paper, we prefer direct methods which work solely with dialogues without appeal to outside methods.

Defining ‘logic’ in this way highlights the importance of the composition problem: Solving the composition problem for a given set L of formulas is a prerequisite for declaring L a logic. Further, this definition of ‘logic’ helps bring to light our fundamental question: When are dialogical “logics” really logics? Not all combinations of structural rules with the standard particle rules result in a logic: There are rulesets where a negative answer to the composition problem can easily be given. For example, let CL' be CL with D10 modified such that \mathbf{P} is now allowed to also assert atoms in defense of disjunctions. Then, $\vDash_{CL'} p \vee \neg p$ and $\vDash_{CL'} (p \vee \neg p) \rightarrow p$, but $\not\vdash_{CL'} p$. The set of formulas CL' for which \mathbf{P} has a winning CL' -strategy is therefore *not* a logic. A less trivial example is the ruleset composed of D10 + D12 + D13; the set of formulas for which \mathbf{P} has a winning strategy according to this ruleset is not closed under *modus ponens*, as both $\neg\neg\varphi \rightarrow \varphi$ and $(\neg\neg\varphi \rightarrow \varphi) \rightarrow (\varphi \vee \neg\varphi)$ are valid, but $\varphi \vee \neg\varphi$ is not. (The same is true if one adds E to the set) [40, Lemma 3.4].

Thus, even if we have justified a set S of structural rules without appealing to some already pre-defined logic that we wish to characterize dialogically, the problem of showing that the set of S -valid sentences form a logic still remains. In particular, in order to justify our calling N a logic, we must show that it meets the requirements of Definition 13. In the next section, we give a positive answer to the strategy composition problem for N , thus proving that the solution to the more general composition problem is positive, hence that it is closed under *modus ponens*, and in §5 we demonstrate closure under non-trivial substitutions.

4. A positive solution to the composition problem for N

In this section we prove the main result of the paper, namely, a positive solution to the composition problem for N. We begin by proving some properties of winning N-strategies.

LEMMA 1. *Every branch for an N-dialogue tree that contains a defensive move by O either terminates at an O-move, or is infinite.*

Proof. If a branch of an N-dialogue tree (i.e., a maximal totally ordered set of nodes of the tree that contains the root) contains such a node a but does not terminate at an O-node, then P has a response to some previous assertion of O. But in this case, O can respond to P's move by repeating the earlier defense of a that occurs in the branch; note that D13 rules out only repeated O-attacks, not repeated O-defenses. Thus branches containing a defensive move for O that do not end with an O-move are infinite.⁹ \square

Corollary 1. No N-winning strategy contains a branch where O defends.

Proof. If s were an N-winning strategy with a branch that contains an defensive O-node a , then, by Lemma 1, every branch of s containing a either terminates at an O-move or is infinite. But since s is a winning strategy, there can be no branches of s that terminate at an O-move, nor can there be any infinite branches. \square

The corollary implies that when every branch of the N-dialogue tree for a formula φ contains a defensive move by O, then φ is N-invalid. The converse, interestingly, fails: In the N-dialogue tree for $(p \rightarrow (\neg q \vee \neg r)) \rightarrow ((\neg p \rightarrow \neg q) \vee (\neg p \rightarrow \neg r))$, which is N-invalid, O never defends in any branch. A consequence of these results is that we can understand N as incorporating (implicitly) Krabbe's rule D4, ∞, ∞ , which allows both O and P to defend against an attack as many times as they like [17, p. 304].

LEMMA 2. *No branch of a winning N-strategy contains an unattacked assertive P-move.*

Proof. There is exactly one fewer O move than there are P-moves in any branch of a winning N-strategy, and the final move is by P. By Corollary 1, all O moves are attacks. If there were an assertive P-move that is unattacked by O, then by the pigeonhole principle there would be a P-move that gets attacked more than once. But this is ruled out by D13. \square

LEMMA 3. *In every branch of a winning N-strategy, if \mathbf{P} makes an assertive move, then \mathbf{O} must respond immediately to this move by attacking \mathbf{P} 's assertion.*

Proof. We do not need to consider defensive moves by \mathbf{O} , by Corollary 1. If \mathbf{O} does not *immediately* attack a \mathbf{P} -assertion when it is made, then \mathbf{O} must be attacking some earlier \mathbf{P} -assertion (note that, by the particle rules, only assertive moves may be attacked). But this violates D13. (We are implicitly employing a proof by induction on the length of the branch.) \square

Thus, although Rule E is missing from \mathbf{N} , some of E's effects are present.

LEMMA 4 (Weakening). *If $\vDash_{\mathbf{N}} \psi$, then $\vDash_{\mathbf{N}} \varphi \rightarrow \psi$, for all formulas φ .*

Proof. Let s_ψ be an N-winning strategy for ψ . The N-dialogue tree $T_{\varphi \rightarrow \psi}$ for $\varphi \rightarrow \psi$ begins with \mathbf{P} 's assertion of $\varphi \rightarrow \psi$, followed by \mathbf{O} 's attacking assertion φ . These first two nodes of $T_{\varphi \rightarrow \psi}$ themselves form a two-element chain, c . Carry out the following modification on s_ψ :

- The root node r of s_ψ is an assertion by \mathbf{P} of ψ , but it is neither an attack or a defense, and it refers to no prior assertion. Change r so that it is now an assertion by \mathbf{P} of ψ , but it is now to be understood as an attack against move 1 (which, in the tree $s_{\varphi \rightarrow \psi}$ that we eventually define, will be \mathbf{O} 's attacking assertion φ against \mathbf{P} 's assertion of $\varphi \rightarrow \psi$);
- Every non-root node of s_ψ refers to some previous assertion number k ; change this to $k + 2$.

Call the result of this modification s'_ψ . Let $s_{\varphi \rightarrow \psi}$ be the result of grafting s'_ψ to the end of c . Claim: $s_{\varphi \rightarrow \psi}$ is an N-winning strategy for \mathbf{P} for $\varphi \rightarrow \psi$. That $s_{\varphi \rightarrow \psi}$ is a subtree of the full N-dialogue tree $T_{\varphi \rightarrow \psi}$ with the same root should be clear (the surgery we carried out on s was intended to ensure that). The more interesting possibility that needs to be ruled out is that in $T_{\varphi \rightarrow \psi}$, \mathbf{O} can respond in more ways than were possible in T_ψ . But this cannot be: D13 is still in force, so that \mathbf{O} can attack \mathbf{P} 's assertions at most once. This implies that \mathbf{O} 's attack against the initial assertion $\varphi \rightarrow \psi$ cannot be repeated, so that any attack by \mathbf{O} must be against some assertion by \mathbf{P} made at some depth ≥ 2 in $T_{\varphi \rightarrow \psi}$; we need not consider defensive moves by \mathbf{O} because of Corollary 1. It remains only to show that every branch of $s_{\varphi \rightarrow \psi}$ is finite and terminates with a \mathbf{P} -move. But this is so because s_ψ has the same property. \square

THEOREM 2 (Characterization of implication). *If $\vDash_{\mathbf{N}} \varphi \rightarrow \psi$ then $\varphi \rightarrow \psi$ satisfies one of the following three conditions:*

1. φ is atomic.
2. φ is negated.
3. $\models_{\mathbf{N}} \psi$.

Proof. Case (3) is simply a restatement of Lemma 4. Suppose now that φ is not atomic and ψ is not an N-validity. We have three cases:

- If φ is an implication $\alpha \rightarrow \beta$, then the N-dialogue tree opens with \mathbf{O} attacking the initial statement by asserting $\alpha \rightarrow \beta$. In any N-winning strategy for $(\alpha \rightarrow \beta) \rightarrow \psi$, \mathbf{P} cannot attack \mathbf{O} 's assertion of $\alpha \rightarrow \beta$, because this leaves open the possibility of a defense by \mathbf{O} , contradicting Corollary 1. Thus, any winning strategy s must choose, for \mathbf{P} 's response to \mathbf{O} 's initial attack, to defend by asserting the consequent ψ of the entire formula, and no branch of s can attack the antecedent implication $\varphi \rightarrow \psi$. By renumbering the reference labels for nodes of s below the \mathbf{P} 's assertion of ψ in the obvious way (renumber k to $k - 2$), we obtain a winning strategy for ψ , contradicting our assumption.
- Likewise, φ cannot be a disjunction, nor could it be a conjunction, for similar reasons: In any N-winning strategy s for $(\alpha \vee \beta) \rightarrow \psi$ (or for $(\alpha \wedge \beta) \rightarrow \psi$), \mathbf{P} never attacks $\alpha \vee \beta$ (respectively, $\alpha \wedge \beta$), so we can recover from s a winning strategy for ψ , contradicting our assumption.

The only remaining possibility is that φ is a negation. □

Remark. To illustrate cases (1) and (2) of this classification of valid implications, consider $p \rightarrow p$ and $\neg p \rightarrow \neg p$. Illustrating (2), we have the more interesting validities $\neg(\varphi \wedge \psi) \rightarrow (\neg\varphi \vee \neg\psi)$ and $\neg(\varphi \vee \psi) \rightarrow (\neg\varphi \wedge \neg\psi)$ (the only directions of De Morgan's laws that are N-valid).

Remark. These conditions are not sufficient: The implicational version of modus ponens, $p \rightarrow ((p \rightarrow q) \rightarrow q)$, has an atomic antecedent, but is (surprisingly) N-invalid.

We can prove a positive solution to the composition problem for \mathbf{N} from the Characterization Theorem and a few simple lemmas:

LEMMA 5. *No atomic formula is N-valid.*

Proof. By Rule D10, the set of N-dialogue trees for an atomic formula p is empty, and hence \mathbf{P} has no winning strategy for any atomic formula. □

Corollary 2. \mathbf{N} is consistent (there is at least one formula φ such that φ is not \mathbf{N} -valid.)

Consistency is closely related to the composition problem; as Krabbe points out:

If we assume that \perp is a statement for which \mathbf{P} has no winning strategy and that $\neg A$ is dialogically equivalent to $A \rightarrow \perp$, consistency is a special case of a metatheorem to the effect that, for all statements A and B , if \mathbf{P} has winning strategies for both $A \rightarrow B$ and for A , he has one for B [19, p. 685].

Thus, by Corollary 2, we know that the composition problem for \mathbf{N} is not trivially solved.

THEOREM 3. *If $\models_{\mathbf{N}} \neg\varphi$, then $\varphi := \neg\psi$ for some ψ for which $\models_{\mathbf{N}} \psi$.*

Proof. By cases:

- φ cannot be atomic, since no negated atoms are \mathbf{N} -valid, by D10 and the particle rule for negation.
- φ cannot be a disjunction $\alpha \vee \beta$ because once \mathbf{O} attacks the negated disjunction by asserting $\alpha \vee \beta$, the only response for \mathbf{P} is to attack the disjunction; \mathbf{O} can (indeed, must) defend by selecting either the left or the right disjunct, so by Corollary 1 no winning strategy exists from this unique initial segment of the \mathbf{N} -dialogue tree for $\neg(\alpha \vee \beta)$.
- Likewise, φ cannot be an implication or a conjunction.

Thus $\varphi = \neg\psi$ for some formula ψ . A winning \mathbf{N} -strategy s_ψ for \mathbf{P} for ψ can be obtained by from a winning strategy $s_{\neg\neg\psi}$ for $\neg\neg\psi$ and noting that, by the particle rule for negation, the winning strategy for $\neg\neg\psi$ begins with a unique initial segment of length two, after which \mathbf{P} asserts ψ , attacking \mathbf{O} 's assertion of $\neg\psi$. Simply remove the root and its unique successor from $s_{\neg\neg\psi}$, declare that \mathbf{P} 's assertion at the new root is neither an attack nor a response, and is a response to no move of \mathbf{O} ; then renumber the reference labels k on all nodes of $s_{\neg\neg\psi}$ by $k - 2$. This renumbering is coherent because neither \mathbf{P} nor \mathbf{O} can attack or respond to moves 0 and 1, by the particle rule for negation and D13, so all reference labels are at least 2. \square

THEOREM 4 (Composition). *If $\models_{\mathbf{N}} \varphi$ and $\models_{\mathbf{N}} \varphi \rightarrow \psi$, then $\models_{\mathbf{N}} \psi$.*

Proof. By the Characterization Theorem 2, given that $\models_{\mathbf{N}} \varphi \rightarrow \psi$, it follows that either

1. φ is atomic,

2. φ is negated, or
3. ψ is N-valid.

Case (3) is the desired conclusion. Case (1) is impossible, in light of the assumption that $\vDash_{\mathbf{N}} \varphi$, by Lemma 5, so the desired conclusion follows vacuously.

It remains to treat case (2). By Theorem 3, from $\vDash_{\mathbf{N}} \varphi$, it follows that $\varphi = \neg\neg\chi$ for some formula χ . The beginning of $T_{\varphi \rightarrow \psi}$ is as follows:

$$\left| \begin{array}{l|l} 0 & \mathbf{P} \\ 1 & \mathbf{O} \end{array} \right| \begin{array}{l} \neg\neg\chi \rightarrow \psi \\ \neg\neg\chi \end{array} \left| \begin{array}{l} (initial\ move) \\ [A, 0] \end{array} \right|$$

Since these are the first two steps of an N-winning strategy for \mathbf{P} , the game does not end here with \mathbf{O} . If, in any branch of s , \mathbf{P} chooses to attack move 1 by asserting $\neg\chi$ as an attack on \mathbf{O} 's assertion of $\neg\neg\chi$, then the dialogue would proceed as follows:

$$\left| \begin{array}{c} \vdots \\ k \\ k+1 \end{array} \right| \left| \begin{array}{c} \mathbf{P} \\ \mathbf{O} \end{array} \right| \left| \begin{array}{c} \neg\chi \\ \chi \end{array} \right| \left| \begin{array}{c} [A, 1] \\ [A, k] \end{array} \right|$$

Such a branch ends with \mathbf{O} , so if there were an N-winning strategy for \mathbf{P} that begins in this way, then \mathbf{P} must have a response. \mathbf{P} cannot attack \mathbf{O} 's assertion of χ at any further point of any branch that begins this way, by Corollary 1. Thus, \mathbf{P} must eventually defend against the attack of move 1 by asserting ψ . We can conclude that \mathbf{P} must actually possess a winning strategy for ψ that can be obtained from s by simply removing all copies of the two-step piece where \mathbf{P} attacks $\neg\neg\chi$. Note that χ cannot be atomic, by Theorem 3, since we are assuming $\vDash_{\mathbf{N}} \neg\neg\chi$. Thus, deleting all these copies of the two-step exchange cannot affect rule D10. Rule D13 is preserved because if \mathbf{O} attacks a \mathbf{P} -statement in the diminished game then the same \mathbf{P} -assertion would likewise be attacked multiple times in the original game. \square

We have thus shown, via semantic means only—that is, our proof does not have to detour through any proof theory; instead, it uses only the semantic infrastructure provided by winning strategies in dialogue games—a positive solution to the composition problem for \mathbf{N} ; this is, to our knowledge, the only solution via purely dialogical means for a composition problem to date.

In order to conclude that \mathbf{N} is a logic, we must show that there exist non-trivial validity-preserving substitutions. We do so in the next section, where we discuss properties of \mathbf{N} .

5. Properties of N

As with the dialogical connexive and relevance logics discussed above, **N** is not closed under unrestricted uniform substitution.¹⁰ Consider, for example, the **N**-validity $p \rightarrow \neg\neg p$ under the substitution of $p \wedge p$ for p : The result of the substitution is **N**-invalid, because the implication no longer meets any of the requirements in the Characterization Theorem 2. The fact that uniform substitution of, e.g., $p \wedge p$ for p in an **N**-valid formula φ is not validity preserving points to a curious type of “resource sensitivity” in the logic—what is valid with some minimal amount of information may fail to remain valid when more information are provided¹¹—though it is not clear exactly how this should be understood, it does show that **N** is a type of substructural logic [33]. The resource sensitivity exhibited by **N** differs from the “limited use” resource sensitivity of linear logic. In linear logic, one is restricted by the number of times a formula can be “used”, whereas **N** is sensitive to the syntactic shape of the formula.

Another counterexample to uniform substitution is the failure of the passage from the **N**-validity $p \vee \neg p$ to an instance $(p \wedge p) \vee \neg(p \wedge p)$. The latter is **N**-invalid: After **O**’s initial attack, in all branches of the **N**-dialogue tree **P** either refrains from asserting $\neg(p \wedge p)$ or asserts it at some move. Branches where **P** asserts $\neg(p \wedge p)$ do not lead to a win for **P** because, after **P**’s assertion of the negation, **O** must defend by asserting the conjunction. This leaves **P** with two options: To attack **O**’s conjunction (and thus fail to win, by Corollary 1), and simply restart the game by defending against the initial attack (so that our analysis of the possible branches recurs, and **P** does not win). Branches in which **P** refrains from asserting $\neg(p \wedge p)$ are infinite because the atomic formula p is never asserted by **O**, so the only way to play for **P** is to infinitely repeat the initial defense against the initial attack, which of course does not lead to a win for **P**.

Despite the failure of unrestricted uniform substitution for **N**, not all is lost. Some types of substitution which are **N**-validity preserving:

LEMMA 6. (1) *Alphabetic renaming of atoms*, (2) *negating atoms*, and (3) *substituting N-validities for atoms are N-validity preserving*.

Proof.

1. A substitution σ of atoms for atoms, suitably extended to formulas, dialogue games, and dialogue trees, clearly sends a dialogue tree $T_{\mathbf{N},\varphi}$ for a formula φ to an isomorphic dialogue tree $T_{\mathbf{N},\sigma(\varphi)}$, since all the particle and structural rules are neutral with respect to atoms (that is, they do not mention any particular atoms).

2. We prove just the cases of (a) singly negating and (b) double negating atoms; higher numbers are provable in an analogous fashion. Further, it suffices to consider singleton substitutions that simply send one atom p to $\neg \dots \neg p$; substitutions which map multiple atoms p, q, \dots are handled simply by decomposing the substitutions into “singleton” substitutions.

(a) Suppose that $\models_N \varphi$, and let s be an N-winning strategy for φ ; we shall construct a winning strategy s^* for $\varphi[p/\neg p]$. First, substitute $\neg p$ for p everywhere in s ; call the resulting tree $s[p/\neg p]$. This tree is isomorphic to s ; let f be an isomorphism, which allows us to speak of branch b in s and its corresponding branch $f(b)$ in $s[p/\neg p]$. The tree $s[p/\neg p]$ is not, in general, an N-winning strategy for $\varphi[p/\neg p]$. We shall modify it so that it becomes one.

There are two kinds of branches of s : those where p is asserted by \mathbf{P} , and those where \mathbf{P} does not assert p . If \mathbf{P} does not assert p in branch b of s , then $f(b)$ is an N-dialogue for $\varphi[p/\neg p]$ won by \mathbf{P} .

Suppose \mathbf{P} has asserted p in branch b of s ; by D10 this atom was asserted first by \mathbf{O} . By Lemma 9, we know that \mathbf{P} 's atomic assertion is the leaf of b . In the corresponding branch $f(b)$, \mathbf{O} instead asserts $\neg p$, and $f(b)$ terminates with \mathbf{P} 's assertion of $\neg p$. The branch $f(b)$ is an N-dialogue, but it is not (yet) won by \mathbf{P} , because there is a legal continuation of the game: \mathbf{O} can attack \mathbf{P} 's assertion of $\neg p$ by asserting p . Extend $f(b)$ with this attack, and then let \mathbf{P} counterattack by asserting p . We see that $f(b)$, extended by these two moves, is an N-dialogue that has ended with a win for \mathbf{P} .

It is clear that s^* is an N-winning strategy for $\varphi[p/\neg p]$: All new moves for \mathbf{O} that have become possible by our transformation are accounted for with responses by \mathbf{P} .

(b) To show that substitutions sending an atom p to its double negation $\neg\neg p$ are also validity preserving we adapt the previous argument. As before, there are two kinds of branches of s : those where p is asserted by \mathbf{P} , and those where \mathbf{P} does not assert p . If \mathbf{P} does not assert p in branch b of s , then $f(b)$ is an N-dialogue for $\varphi[p/\neg p]$ won by \mathbf{P} .

Suppose \mathbf{P} has asserted p in branch b of s ; by D10 this atom was asserted first by \mathbf{O} . By Lemma 9, we know that \mathbf{P} 's atomic assertion is the leaf of b . In the corresponding branch $f(b)$, \mathbf{O} instead asserts $\neg\neg p$, and $f(b)$ terminates with \mathbf{P} 's assertion of $\neg\neg p$. The branch $f(b)$ is an N-dialogue, but it is not (yet) won by \mathbf{P} , because there is a legal continuation of the game: \mathbf{O} can attack \mathbf{P} 's assertion of $\neg\neg p$ by asserting $\neg p$. Extend $f(b)$ with this attack. We cannot

conclude that the game is over, because it may be that **P** cannot yet assert p . Note, however, that in b , we find that **O** first asserted p ; in $f(b)$, **O** has asserted $\neg\neg p$. Extend $f(b)$ with an attack by **P** against **O**'s assertion of $\neg\neg p$. **O** must respond immediately to this attack, so **O** asserts p . **P** is now in a position to respond to **O**'s assertion of $\neg p$. We see that $f(b)$, extended by these four moves, is an N-dialogue that has ended with a win for **P**.

The arguments just given for the case $p \mapsto \neg p$ and $p \mapsto \neg\neg p$ can be adapted to deal with substitutions that send an atom to any number of negations of itself. For the even case, we adapt the proof above for $p \mapsto \neg\neg p$, and for the odd case we adapt the proof above for $p \mapsto \neg p$.

3. Let s be an N-winning strategy for φ , p be an atom in φ , α be an N-validity, and t be an N-winning strategy for α . By Lemma 9, we know that every branch b of s terminates with the first occurrence by **P** of some atom, that is, b ends with a **P**-assertion of an atom, and there is no atom that **P** asserts earlier in b . We shall construct an N-winning strategy s^* for $\varphi[p/\alpha]$.

First, substitute α for p everywhere in s ; call the resulting tree $s[p/\alpha]$. This tree is isomorphic to s ; let f be the isomorphism, which allows us to speak of branch b in s and its corresponding branch $f(b)$ in $s[p/\alpha]$. However, $s[p/\alpha]$ is not, in general, an N-winning strategy for $\varphi[p/\alpha]$. We shall modify it so that it becomes one.

As in the previous case, there are two kinds of branches of s : those where p is asserted by **P**, and those where **P** does not assert p . If **P** does not assert p in branch b of s , then $f(b) = b$. Suppose **P** has asserted p in branch b of s . By Lemma 9, this atomic assertion is a leaf in s . Branch $f(b)$, however, is not maximal, since, by Lemma 5, its leaf is a complex assertion by **P**. However, we can make $f(b)$ maximal by attaching the strategy t for α (with labels suitably adjusted) to the end of $f(b)$.¹² Let s^* be the result of this process.

Claim. This strategy s^* is an N-winning strategy for $\varphi[p/\alpha]$.

Proof. We have to show that:

- a) Each branch of s^* is an N-dialogue won by **P**.

Since we started with an N-winning strategy s , all branches b in s are N-dialogues won by **P**. Thus, in a branch of s where **P** does not assert p , the corresponding branch of s^* is also a N-dialogue won by **P**. In a branch where **P** has asserted p , the

corresponding subtree of s^* is such that **P** wins every branch, since t was a winning strategy for α .

- b) For every node u of s^* corresponding to a move by **P**, the set of immediate children of u is exactly the set of all N-permissible moves for **O**.

This follows straightforwardly from the fact that s and t are N-winning strategies.

□
□

A straightforward consequence of this, especially case (2), is the following:

LEMMA 7. *N-validity is preserved by double negating an arbitrary subformula.*

Proof. Suppose that s is an N-winning strategy for φ , and let φ^* be a replacement in φ of a subformula, ψ , by its double negation $\neg\neg\psi$. We shall describe an N-winning strategy s^* for φ^* . If ψ never occurs in any branch of s , then $s = s^*$. Otherwise, consider a branch b of s where φ is in fact asserted. In s^* , the node corresponding to the assertion of ψ is an assertion of $\neg\neg\psi$; suppose this occurs at move k of b . Modify s by inserting the following:

$$\left| \begin{array}{c|c|c} k & \mathbf{P/O} & \neg\neg\psi \\ k+1 & \mathbf{O/P} & \neg\psi \\ k+2 & \mathbf{P/O} & \psi \end{array} \middle| \begin{array}{c} [A/D, x] \\ [A, k] \\ [A, k+1] \end{array} \right|$$

From $k+3$, **P** can continue playing by s (with the move references adjusted by 2, to account for the insertion of two new moves in b). The result of this insertion is s^* , and it is a winning N-strategy for **P** since we have considered all of the new possible moves of **O**. □

Corollary 3. N-validity is preserved by negative translations that simply add double negations, such as Kolmogorov's and Kuroda's [39, p. 58]¹³, and by negating arbitrary subformulas an even number of times.

The fact that N validates some, but not all, types of uniform substitution raises a question concerning its substitution core.

Definition 14. The *substitution core* of a logic \mathcal{L} is the set of schematic validities of \mathcal{L} .

In logics which satisfy uniform substitution, the set of valid formulas and the set of schematically valid formulas coincide; but when uniform substitution is dropped, these two sets will diverge.

Question. What is the substitution core of N?

Having established that N is a logic, we next say something about what type of logic it is, and how it fits into the scheme of known propositional logics. In Theorem 2 we gave a characterization of N-valid implications, and in Theorem 3 one of N-valid negations. We also know by Lemma 5 that no atom is N-valid. We now give similar results for the remaining connectives.

As in intuitionistic and classical logic, validity of a conjunction transmits to the conjuncts:

THEOREM 5. $\vDash_N \varphi \wedge \psi$ iff $\vDash_N \varphi$ and $\vDash_N \psi$.

Proof. Straightforward. □

Likewise, weakening by disjunction is valid:

THEOREM 6. *If $\vDash_N \varphi$ or $\vDash_N \psi$, then $\vDash_N \varphi \vee \psi$.*

Proof. Assume $\vDash_N \varphi$. Then **P** has a winning strategy for φ . We now show that he also has a winning strategy for $\varphi \vee \psi$. The N-dialogue tree $T_{\varphi \vee \psi}$ for $\varphi \vee \psi$ begins with **P**'s assertion of $\varphi \vee \psi$ followed by **O**'s attack ? on $\varphi \vee \psi$. In round three, let **P** defend this attack by asserting φ . He is now able to use his winning strategy for φ ; since **O** had no other choice of move at round two, this constitutes a winning strategy for $\varphi \vee \psi$.

The case for $\vDash_N \psi$ is symmetric. □

In IL, the converse of Theorem 6 holds, but it does not hold in either CL or N. For example, $p \vee \neg p$ is valid in N, but neither p nor $\neg p$ is. We can, however, characterize N-valid disjunctions in a way similar to the way CL-valid disjunctions can be characterized:

THEOREM 7. *If $\vDash_N \varphi \vee \psi$, then there is an instance of an atom that does not appear in the antecedent of a conditional which is a subformula of $\varphi \vee \psi$, and there is either some instance of that atom in the antecedent of a conditional with the same parity of negations or some instance of that atom not in the antecedent of a conditional with the opposite parity of negations.*

We prove this with the help of the following lemmas:

LEMMA 8. *If an N-dialogue game ends with a win for \mathbf{P} , then it ends with \mathbf{P} 's assertion of an atom.*

Proof. If an N-dialogue d ends with a win for \mathbf{P} , then \mathbf{P} has made the final assertion of the dialogue and there is no further legal move that \mathbf{O} could make. Suppose that d does not end with the assertion by \mathbf{P} of an atom. By the particle rules, there are two cases:

- If \mathbf{P} 's final move is a defense, then, by the particle rules, \mathbf{P} has asserted a formula. By hypothesis, the formula is not an atom, and hence \mathbf{O} may attack this formula, contradicting the hypothesis that \mathbf{O} cannot legally continue the game.
- If \mathbf{P} 's final move is an attack, there are two subcases:
 - \mathbf{P} asserts a formula. By assumption, the formula is a complex assertion which has not yet been attacked, and hence \mathbf{O} may attack this formula, contradicting the hypothesis that \mathbf{O} cannot legally continue the game.
 - \mathbf{P} makes a symbolic attack. Contra assumption, \mathbf{O} can legally defend against this attack; indeed, \mathbf{O} can always defend against any symbolic attack by \mathbf{P} .

□

Corollary 4. If \mathbf{P} wins an N-dialogue, then there is an atom that \mathbf{P} asserts in the game.

LEMMA 9. *Every branch of every N-winning strategy for φ terminates with the first assertion by \mathbf{P} of an atom.*

Proof. By Corollary 1, \mathbf{O} never defends in a branch b of a winning strategy, so all his moves are attacks. By Lemma 3, \mathbf{O} must immediately attack \mathbf{P} 's assertions. Thus, if \mathbf{P} asserts an atom in b , there must be some attack by \mathbf{O} against this; but such a move is impossible, by the particle rules. □

Every assertion of a formula in an N-dialogue is, by the particle rules, the assertion of a particular subformula of the initial formula of the dialogue. Given the assertion of a formula at some step in a dialogue, one can trace which particular subformula of the initial formula is involved. It can happen that identical formulas are asserted in a game (consider, for example, a game for $p \rightarrow p$), but have different “origins” in virtue of being distinct subformulas of the initial formula.

LEMMA 10. *In any N-dialogue for φ , \mathbf{P} and \mathbf{O} never assert the same occurrence of the same subformulas of φ .*

The main idea of the proof is that the particle rules ensure that the “ancestry” of every formula appearing in an N-dialogue is unique.

Proof. In any N-dialogue d for φ , it is the case that (1) only **P** asserts φ , and (2) if there is an occurrence of a subformula α of φ asserted in d by both **P** and **O**, then there is a (proper) superformula α^* of α such that both **P** and **O** assert this occurrence of α^* in d . Together these imply that there cannot be an occurrence of a subformula of φ asserted by both **P** and **O**, because by repeated applications of (2) we would find successively (proper) superformulas of this subformula, leading eventually to the whole initial formula α ; but by (1), we cannot have both **P** and **O** asserting the initial formula.

To prove (1), note that every move of d after the first is an attack or a defense; by the particle rules, either a symbolic attack is made, or a proper subformula is asserted in attack or defense.

To prove (2), suppose that **P** and **O** both assert the same occurrence of a proper subformula α of the initial formula φ . Since α is a proper subformula of φ , there is a unique immediate superformula α^* of α that is also a (possibly improper) subformula of φ . There are several possibilities, depending on the main connective of α^* , but all can be treated in a similar way. We treat only the conjunction case.

Suppose that α is the left conjunct of $\alpha^* := \alpha \wedge \alpha'$. By the particle rules, the only way for α to occur in the game is if $\alpha \wedge \alpha'$ also occurs. Since players cannot respond to their own statements, it must be the case that both **P** and **O** have asserted α^* . The case where α is the right conjunct or where α^* is a disjunction, negation, or implication are treated similarly.

Since at no point in this proof did we appeal to anything specific about N, the following corollary is immediate:

Corollary 5. For any ruleset S, in any S-dialogue for φ , **P** and **O** never assert the same occurrence of the same subformulas of φ .

We are now in a position to prove Theorem 7.

Proof. Assume $\vDash_N \varphi \vee \psi$. Let s be a winning N-strategy for $\varphi \vee \psi$, and b a branch in s . By Lemma 9, b ends with the assertion by **P** of an atom, call it p . By Lemma 10, we know that p occurs more than once in $\varphi \vee \psi$. By D10, **P** will only assert p if **O** has already asserted it. Since s is a winning strategy, by Corollary 1, **O** will only assert an atom as an attack on some formula, and not as a defense, since **O** makes no defensive moves in b , and the formula being attacked must be either an implication or a negation. Thus, either the occurrence of p that **O**

Table II. Some N-validities

$p \vee \neg p$	$\neg p \vee \neg\neg p$
$(p \rightarrow q) \vee (p \rightarrow \neg q)$	$(p \rightarrow q) \vee (q \rightarrow p)$
$\neg\neg p \rightarrow p$	$p \rightarrow \neg\neg p$
$p \rightarrow (p \vee q)$	$p \rightarrow (p \wedge p)$
$\neg p \rightarrow (p \rightarrow q)$	$\neg(p \vee \neg p) \rightarrow q$

asserts is a subformula of a negated atom with an opposite number of negations to the occurrence that \mathbf{P} asserts, or p is in the antecedent of a conditional. In the first case, we are done. In the second, we must show that \mathbf{P} 's assertion of p is also not in the antecedent of a conditional. But if p is in the antecedent of a conditional, either it is a subformula of the antecedent or it is identical to the antecedent. In either case, \mathbf{P} will only assert p in the course of an attack on the conditional, either by asserting a complex antecedent containing p or by asserting p itself, and such a move can be defended by \mathbf{O} , violating Corollary 1, and the assumption that b is a branch in a winning strategy. \square

We give examples of N-valid formulas in Table II. More generally, we know that

THEOREM 8. $\mathbf{N} \subset \mathbf{CL}$.

Proof. Every D10 + D13-strategy is also a D10 + D13 + E-strategy, by Lemma 3. That the inclusion is strict follows from the fact that $\not\models_{\mathbf{N}} (((p \rightarrow q) \rightarrow p) \rightarrow p)$ (Peirce's law), which is classically valid. \square

As a corollary, \mathbf{N} is not a connexive logic. We also know that:

LEMMA 11. $\mathbf{N} \not\subseteq \mathbf{IL}$ and $\mathbf{IL} \not\subseteq \mathbf{N}$.

Proof. For the first claim, $\models_{\mathbf{N}} p \vee \neg p$. For the second claim, $\models_{\mathbf{IL}} (\neg p \vee \neg q) \rightarrow \neg(p \wedge q)$, which, by Theorem 2 is not N-valid, since $\not\models_{\mathbf{N}} \neg(p \wedge q)$. \square

Since \mathbf{N} is neither sub-intuitionistic nor super-intuitionistic, but is sub-classical, it lies in an interesting and as yet under-investigated part of the lattice of propositional logics.

\mathbf{N} shares many characteristics with known sub-classical propositional logics, though it does not completely align with any of them. It turns out that although $\models_{\mathbf{N}} \varphi \wedge \psi$ iff $\models_{\mathbf{N}} \varphi$ and $\models_{\mathbf{N}} \psi$, and $\models_{\mathbf{N}} \varphi \rightarrow \psi$ implies $\models_{\mathbf{N}} \neg\psi \rightarrow \neg\varphi$, these results and others like them do not hold when formulated as object-language implications. For example, conjunction elimination, considered as a formula— $p \wedge q \rightarrow p$ —is not N-valid, and,

even more surprisingly, given that \mathbf{N} is closed under *modus ponens*, neither the conjunctive $((p \wedge (p \rightarrow q)) \rightarrow q)$ nor the implicational $(p \rightarrow ((p \rightarrow q) \rightarrow q))$ versions of *modus ponens* is \mathbf{N} -valid. Thus, the fact that double negation introduction and elimination are both valid is noteworthy.

Although \mathbf{N} does not validate unrestricted uniform substitution, it is nevertheless consistent with it.

Definition 15. Let \mathbf{N}^* be the smallest set of formulas extending \mathbf{N} that is closed under unrestricted uniform substitution.

THEOREM 9. \mathbf{N}^* is consistent.

The proof exploits the fact that \mathbf{N} is subclassical.

Proof. Suppose that there were a formula φ such that both $\varphi \in \mathbf{N}^*$ and $\neg\varphi \in \mathbf{N}^*$. Because \mathbf{N} is consistent, we know that it's not the case that $\varphi, \neg\varphi \in \mathbf{N}$. Thus, there are formulas α and β in \mathbf{N} such that φ is obtained by an application of uniform substitution from α and likewise $\neg\varphi$ is obtained by an application of uniform substitution from β . Since $\alpha \in \mathbf{N}$, by Theorem 8, α is a classical tautology; and since classical logic \mathbf{CL} is closed under uniform substitution, φ is a classical tautology, too. Likewise, $\neg\varphi$ is a classical tautology. But both φ and $\neg\varphi$ can't be tautologies, because \mathbf{CL} is consistent.¹⁴ \square

6. A tableau system for \mathbf{N}

\mathbf{N} , which is characterized by winning strategies under a set of dialogue rules, does not have natural semantics or proof theory. In this section, we give a set of tableau rules which allow us to make a first step at giving an alternate, non-dialogical characterization of the logic.

Tableaux are an efficient proof procedure for propositional and predicate logics, and have a natural connection to dialogues via extensive games [5, 30]. The tableaux we construct for \mathbf{N} are adapted from the standard analytic \mathbf{CL} -tableaux for signed formulas [7, 35].¹⁵

Definition 16. If φ is an unsigned formula, then $\mathbf{P}\varphi$ and $\mathbf{O}\varphi$ are signed formulas.

Definition 17. An \mathbf{N} -tableau for $\mathbf{X}\varphi$ is defined inductively (following [7, Def. 6, p. 51]):

1. $\mathbf{X}\varphi$ is an \mathbf{N} -tableau for $\mathbf{X}\varphi$.

Table III. Tableau rules for N.

$$\begin{array}{l}
\mathbf{O} \neg \frac{\mathbf{O} \neg \varphi}{\mathbf{P} \varphi} \\
\mathbf{O} \rightarrow \frac{\mathbf{P}(\varphi \rightarrow \psi)}{\mathbf{O} \varphi} \\
\mathbf{P} \neg \frac{\mathbf{P} \neg \varphi}{\mathbf{O} \varphi} \\
\mathbf{P} \wedge \frac{\mathbf{P}(\varphi \wedge \psi)}{\mathbf{P} \varphi | \mathbf{P} \psi} \\
\mathbf{P} \vee_L \frac{\mathbf{P}(\varphi \vee \psi)}{\mathbf{P} \varphi} \\
\mathbf{P} \vee_R \frac{\mathbf{P}(\varphi \vee \psi)}{\mathbf{P} \psi} \\
\mathbf{P} \rightarrow \frac{\mathbf{P}(\varphi \rightarrow \psi)}{\mathbf{P} \psi}
\end{array}$$

2. If \mathcal{T} is an N-tableau for $\mathbf{X}\varphi$, and R is one of the rules in Table III, then the extension of \mathcal{T} to $\mathbf{R}(\mathcal{T})$ by this rule is an N-tableau for $\mathbf{X}\varphi$.
3. Nothing else is an N-tableau for $\mathbf{X}\varphi$.

Thus, tableaux are labeled rooted trees. A tableau is called *proper* if no rule is applied more than once in a branch. In what follows, we will only consider proper tableaux.

For any of the tableau rules in Table III, we call the formula above the line the *premise* of the rule, and any formula below the line a *conclusion* of the rule. Note that these rules are to be understood as rule *schemes*. Thus, strictly speaking, there are not eight rules, but countably infinitely many, one for every combination of formulas in the language. For example, for $\varphi \neq \psi$,

$$(1) \frac{\mathbf{O} \neg \varphi}{\mathbf{P} \varphi}$$

and

$$(1) \frac{\mathbf{O} \neg \psi}{\mathbf{P} \psi}$$

are, despite the identity of the label (1) in the two figures, applications of two *different* rules, not two applications of the same rule. We should, properly, index each of the rules above to the formulas φ and ψ appearing in them, but we omit doing so for simplicity's sake.




Figure 1. A simple, \mathbf{O} -preferring tableau for $\mathbf{P}\neg(a \vee b) \rightarrow (\neg a \vee \neg b)$

Definition 18. (Closure (atomic)) A branch of an \mathbf{N} -tableau is said to be *closed* if there is an atom p such that the branch contains both $\mathbf{P}p$ and $\mathbf{O}p$. An \mathbf{N} -tableau is said to be *closed* iff every branch of the tableau is closed.

If a branch is closed because of p , we call p a *closure atom*. Note that a branch need not have a unique closure atom. Without loss of generality we can assume that in the construction of a tableau, all branches close as soon as they can; that is, as soon as we have $\mathbf{O}p$ and $\mathbf{P}p$ in the branch, we no longer expand that branch.

Definition 19. A branch B in a tableau is *saturated* if every rule that can be applied in B has been. In other words, for every node u in B , if a rule R can be applied to u , then there is some descendant of u in B that is obtained by applying R to u .

Definition 20. A tableau \mathcal{T} is *complete* if every branch is either closed or saturated.

Since we are dealing here only with proper tableaux, any branch in a complete tableaux which does not close will have exactly one application of every rule that it's possible to apply.

Definition 21. A tableau \mathcal{T} is *simple* if each node is generated by an application of a rule to the formula of least complexity in the branch, where the complexity of a formula is the number of connectives it has.

Definition 22. A tableau \mathcal{T} is *\mathbf{O} -preferring* if for all branches B of \mathcal{T} and all nodes u of B whose label is a \mathbf{P} -signed implication, the $\mathbf{O} \rightarrow$ rule is applied to u in B before the $\mathbf{P} \rightarrow$ rule.

An example of a simple, \mathbf{O} -preferring closed tableau is given in Figure 1.

Definition 23. A tableau \mathcal{T} is *left-preferring* if for all branches B of \mathcal{T} and all nodes u of B , if the label of u is a \mathbf{P} -signed disjunction, then $\mathbf{P}\vee_L$ rule is applied to u in B before the $\mathbf{P}\vee_R$ rule.

N-tableaux differ from CL-tableaux in three ways. First, we reinterpret (and hence relabel) the signed formulas. Since the origin of **N** is fully dialogical, we do not have semantical notions of truth and falsity at our disposal. Therefore instead of signing our formulas with truth values, we sign them with players in the dialogue game, mapping **O** to **T** and **P** to **F**. Thus, a tableau which corresponds to a winning strategy for **P** in a dialogue game for φ will have $\mathbf{P}\varphi$ as its root, corresponding to an initial assertion of φ by **P**.

Second, the defining characteristic of **N** is that any branch where **O** is allowed to defend is infinite (cf. Lemma 1), and hence no winning strategy for **P** will contain any such branch. There are two ways that we can modify the tableaux rules for **CL** to reflect this property of **N**. One, we can simply remove the rules that would correspond to defensive moves by **O**; the rules modified in this way are given in Table III. Two, we could retain the defensive rules for **O**, but change them so that any defense could be repeated. This is done by copying the defended formula into the defenses, thus ensuring that any time a defensive move by **O** is made, it will still be possible to repeat that move. The tableaux generated by the first set of rules are simpler, and hence these are the rules that we use.

Third, the tableau rules for **CL** are generally stated in a multiple consequent form. Whereas we have separated out the rule for disjunction into two rules $\mathbf{P}\vee_L$ and $\mathbf{P}\vee_R$, the corresponding rule in **CL** is a single rule which extends the tree by two nodes.

7. Winning strategies and N-tableaux

Our goal in this section is to prove the following:

THEOREM 10. *The following are equivalent:*

1. φ is N-valid.
2. **P** has a winning strategy in an N-dialogue for φ .
3. Every N-tableau for $\mathbf{P}\varphi$ is closed.

Proof.

- (1) \Rightarrow (2) and (2) \Rightarrow (1) are trivial; the notion of N-validity is defined in terms of the existence of **P**-winning N-strategies for φ .
- (2) \Rightarrow (3): This is Corollary 7.
- (3) \Rightarrow (2): This is Lemma 15.

□

We show that (2) implies (3) by showing that if there is a closed N-tableau for a formula φ , then every N-tableau for φ is closed. We then define an algorithm that maps winning N-strategies to closed N-tableaux. For showing that (3) implies (2), we give a simple algorithm which converts a closed, simple, **O**-preferring tableau for $\mathbf{P}\varphi$ into a **P**-winning N-strategy for φ .

7.1. CANONICAL N-TABLEAUX

Definition 24. Given a formula φ , we call the unique simple, **O**-preferring, left-preferring N-tableau \mathcal{T} for $\mathbf{P}\varphi$ the *canonical N-tableau for $\mathbf{P}\varphi$* , denoted $C(\varphi)$.

The existence of a canonical tableau for $\mathbf{P}\varphi$ is obvious. We now justify uniqueness. Since the root of $C(\varphi)$ is a single formula, in the construction of a simple tableau from $\mathbf{P}\varphi$, there is always a unique formula of least complexity that still has rules that can be applied to it. Whenever there is a choice of rules to be applied, since $C(\varphi)$ is both **O**-preferring and left-preferring, **O** \rightarrow is applied before **P** \rightarrow , and **P** \vee_L before **P** \vee_R . Hence, at any stage in the construction of $C(\varphi)$, there is a unique formula to be apply a rule to, and a unique rule to be applied. Thus, there is only one way that $C(\varphi)$ can be constructed. (Recall here that we are only dealing with proper tableaux, hence no rule will be applied more than once.)

We now define a reduction relation for N-tableaux, and show that every N-tableau for $\mathbf{P}\varphi$ can be reduced to the canonical N-tableau for $\mathbf{P}\varphi$.

Definition 25. (Reduction relation for N-tableaux) We say that $\mathcal{T} \rightarrow \mathcal{T}'$ if \mathcal{T}' is obtained from \mathcal{T} in the following way: Given a node u of \mathcal{T} , let B be the branch of \mathcal{T} to which u belongs, and let $X\varphi$ be the signed formula belonging to B such that φ is of minimal complexity among all formulas appearing in u above B . If (the formula labeling the node) u is not the premise of an application of an N-tableau rule R to $X\varphi$, then rewrite \mathcal{T} as follows:

- (a) If $X = \mathbf{P}$ and φ is a conjunction, then make u the premise of an application of **P** \wedge . If there are applications below u of **P** \wedge to u , they are to be deleted. Rewrite the tree as follows:
 - Let v_l be an application of **P** \wedge to u below u in the left-hand branch of the new subtree rooted at u , let v_r be the corresponding application of **P** \wedge to u below u in the right-hand branch of the new subtree rooted at u .

- Let v_{l_1} and v_{l_2} be the left and right children, respectively, of the application of $\mathbf{P}\wedge$ to v_l , and likewise let v_{r_1} and v_{r_2} be the children of the application of $\mathbf{P}\wedge$ to v_r .
 - Delete the application of $\mathbf{P}\wedge$ at v_l ; join v_{l_1} to the parent of v_l . Delete the subtree rooted at v_{l_2} entirely.
 - Likewise, delete the application of $\mathbf{P}\wedge$ at v_r ; join v_{r_2} to the parent of v_r . Delete the subtree rooted at v_{r_1} entirely.
- (b) If $X = \mathbf{P}$ and φ is a disjunction, then make u the premise of an application of $\mathbf{P}\vee_L$. If there are any applications of $\mathbf{P}\vee_L$ below u , delete them.
- (c) If $X = \mathbf{P}$ and φ is an implication, then make u the premise of an application of $\mathbf{O}\rightarrow$. If there are any applications of $\mathbf{O}\rightarrow$ below u , delete them.
- (d) If $X = \mathbf{P}$ and φ is a negation, then make u the premise of an application of $\mathbf{P}\neg$. If there are any applications of $\mathbf{P}\neg$ below u , delete them.
- (e) If $X = \mathbf{O}$ and φ is a negation, then make u the premise of an application of $\mathbf{O}\neg$. If there are any applications of $\mathbf{O}\neg$ below u , delete them.

Each of these transformations produces an N-tableau. The last step of each transformation—deleting any duplicate applications of rules just introduced—ensures that the result is a proper N-tableau. Each step “repairs” one node of \mathcal{T} that violates our notion of canonical N-tableau, but though it specifies a way to repair one node, it doesn’t say which node to repair. (Note that if \mathcal{T} is the canonical N-tableau for φ , then none of the above operations apply.)

We define \rightarrow as the reflexive transitive closure of \rightarrow . The following guarantees that there is at least one way to go from an arbitrary N-tableau for φ to $C(\varphi)$ by taking steps along the \rightarrow relation.

Proposition 1. For every formula φ and every N-tableau \mathcal{T} for $\mathbf{P}\varphi$, $\mathcal{T} \rightarrow C(\varphi)$.

Proof. We show this by showing the existence of a finite reduction sequence $\mathcal{T} := \mathcal{T}_0 \rightarrow \mathcal{T}_1 \rightarrow \cdots \rightarrow \mathcal{T}_{n-1} \rightarrow \mathcal{T}_n = C(\varphi)$. Initially, set $\mathcal{T}_0 := \mathcal{T}$. Given \mathcal{T}_k , start at its root and check whether any of the transformations specified in Definition 25 apply there. If so, there is a unique case that applies; apply it, let the result be \mathcal{T}_{k+1} , and continue the process with \mathcal{T}_{k+1} . If none of the transformation rules applies at

the root of \mathcal{T}_k , proceed down to the child (or children) of the root of \mathcal{T}_k and check whether any of the transformation rules applies. (If there are two children of the node of \mathcal{T}_k that we are inspecting, look first at the left child, then at the right.) If a transformation rule applies, then there is a unique one that does apply; apply it, let the result be \mathcal{T}_{k+1} , and restart the process there. Otherwise, continue moving down \mathcal{T}_k , always looking for the shallowest node where one of the transformation rules applies. If none of the rules apply, then we are done.

The reduction sequence is finite because at each stage we always move down at least one level of the tree, and the transformations of Definition 25 cannot grow the N-tableau to which they are applied more than the number of subformulas of φ . \square

Proposition 2. If $\mathcal{T} \rightarrow \mathcal{T}'$, then \mathcal{T} is closed iff \mathcal{T}' is closed.

Proof. Immediate from the definition of \rightarrow : the branches of \mathcal{T}' are introduced by copying branches from \mathcal{T} , so if every branch of \mathcal{T} is closed, then so is every branch of \mathcal{T}' , and vice versa. \square

Corollary 6. If $\mathcal{T} \twoheadrightarrow \mathcal{T}'$, then \mathcal{T} is closed iff \mathcal{T}' is closed.

LEMMA 12. *If there is a closed N-tableau for φ , then every N-tableau for φ is closed.*

Proof. Let \mathcal{T} be a closed N-tableau for φ , and let \mathcal{T}' be an N-tableau for φ . We need to show that \mathcal{T}' is closed. First, we have that $\mathcal{T} \twoheadrightarrow C(\varphi)$ by Proposition 1. Since \mathcal{T} is closed, we have by Corollary 6 that $C(\varphi)$ is likewise closed. Furthermore, $\mathcal{T}' \twoheadrightarrow C(\varphi)$, again by Proposition 1. Corollary 6 then tells us that \mathcal{T}' is closed, as desired. \square

7.2. MAPPING STRATEGIES TO TABLEAUX

We now turn to the problem of showing that winning N-strategies can be mapped to (closed) N-tableaux. To show this, we need the following:

LEMMA 13. *Every branching node u of every winning N-strategy s is a **P**-node whose label is an **O**-signed conjunction $\mathbf{O}\varphi \wedge \psi$ with exactly two children whose labels are the two attacks \wedge_L and \wedge_R against the immediately prior assertion $\varphi \wedge \psi$.*

Proof. Suppose that u is a branching node of a winning N-strategy. Such a node must be a **P**-node, since children of **O**-nodes have exactly one child, by the definition of winning strategy. If there are multiple children of u representing attacks by **O**, then they must be against the immediately prior assertion of **P**, by Lemma 3. And by the particle

rules the only way there can be multiple ways for \mathbf{O} to attack the immediately prior statement by \mathbf{P} is if \mathbf{P} has asserted a conjunction and \mathbf{O} can choose to attack this with either \wedge_L or \wedge_R . \square

LEMMA 14. *Any \mathbf{P} -winning N-strategy for a formula φ can be transformed into a closed N-tableau for $\mathbf{P}\varphi$.*

Proof. We give an algorithm for constructing a tableau from a winning strategy. The algorithm takes as input as strategy s , and the output will be a tree \mathcal{T} which we will prove to be a tableau. In the algorithm we will refer to “the current tableau node” of \mathcal{T} and “the current strategy node” as the node of s that we are operating on. We will indicate how the output tree is constructed, and indicate how to assign tableau rules to the nodes.

Initially, we are given the root of s ; this is the current strategy node. The current tableau node is not yet assigned. The root of s is a \mathbf{P} -node and contains the assertion φ . Make a new node u whose label is $\mathbf{P}\varphi$; this is the root node of \mathcal{T} . Set the current tableau node to be u . Proceed down s as follows:

1. The current strategy node has no children: We are done.
2. The current strategy node has an assertive child: We know that this child is unique; call it c . Make a new child v of the current tableau node; assign the label $X\varphi$ to v where X is the player of c and φ is the formula asserted at this node. Update the current strategy node to c , and update the current tableau node to v . The tableau rule that will label the edge from the current tableau node to the new tableau node v is determined by the player and the stance of the current strategy node c , as well as the formula being attacked or defended. It is determined as follows:
 - a) If c is a \mathbf{P} -node attacking a negation, then the label is $\mathbf{P}\neg$.
 - b) If c is an \mathbf{O} -node attacking a negation, then the label is $\mathbf{O}\neg$.
 - c) If c is a \mathbf{P} -node defending against an attack on a disjunction, then the label is $\mathbf{P}\vee_L$ or $\mathbf{P}\vee_R$ depending on whether \mathbf{P} defends by asserting the left disjunct or the right.
 - d) If c is a \mathbf{P} -node defending against an attack on a conjunction, then the label is $\mathbf{P}\wedge$.
 - e) If c is an \mathbf{O} -node attacking an implication, then the label is $\mathbf{O}\rightarrow$;
 - f) If c is a \mathbf{P} -node defending against an attack on an implication, then the label is $\mathbf{P}\rightarrow$.

Reinitiate the algorithm with c .

3. The current strategy node has a non-assertive child: We know that the child node is an **O**-node, because in a winning N-strategy, **P** will never make any symbolic attacks (by Corollary 1). Furthermore, since s is a winning strategy, any **O**-node will have exactly one **P**-node as successor, and it will be an assertive move. There are two cases:
 - a) The current strategy node has a unique child c , which itself has a unique child c' : Set the new current strategy node to be c' . Create a new tableau node and label it with $\mathbf{P}\varphi$ where φ is the formula asserted by **P** at c' ; make this node the current tableau node. The tableau rule that will label the edge from the current tableau node to the new tableau node is determined as above.
 - b) The current strategy node has two children c_1 and c_2 , each of which has a unique child c'_1 or c'_2 : Set the new current strategy node to be c'_1 , and continue as in the previous step. When the algorithm terminates, set c'_2 to be the new current strategy node, reset the current tableau node to whatever the current tableau was when c'_1 was the current strategy node and continue as in the previous step.

Claim. The output of the algorithm is a tableau \mathcal{T} for $\mathbf{P}\varphi$.

Proof. \mathcal{T} is a tableau because of the direct correspondence between the particle rules of dialogue games and the rules for constructing N-tableaux of Table III. Note that the algorithm ensures that if we add one branch of the rule $\mathbf{P}\wedge$, we will always add the other eventually, since the only time that there is a node in the strategy where **P** defends against an attack on a conjunction is when **O** has attacked a conjunction, and hence the strategy branches.

We give an illustration of this correspondence here. Consider a fragment of an N-dialogue game such that at some step **O** attacks a formula of the form $\varphi \rightarrow \psi$ (previously granted by **P**) by asserting φ and at the next step **P** defends against this attack by asserting ψ . Applying the algorithm to this fragment of the game, we get the following branch:

$$\begin{array}{c}
 \vdots \\
 \mathbf{O} \rightarrow \frac{\mathbf{P}(\varphi \rightarrow \psi)}{\mathbf{P} \rightarrow \frac{\mathbf{O}\varphi}{\mathbf{P}\psi}} \\
 \vdots
 \end{array}$$

It is not difficult to see that this is in fact a fragment of a branch of a tableau. \square

Claim. The tableau \mathcal{T} is closed.

Proof. \mathcal{T} is a closed tableau because of the direct correspondence between \mathbf{P} 's winning conditions for N-dialogues and the definition of N-tableaux closure. \mathbf{P} wins a dialogue if it is \mathbf{O} 's turn but no legal move can be made; thus, whenever we have a \mathbf{P} -winning strategy, we know that in all of its branches \mathbf{O} has exhausted his legal moves. Now suppose we have a winning N-strategy s . Consider an arbitrary branch B of s . In B , \mathbf{P} has answered to all attacks of \mathbf{O} on non-atomic formulas asserted by \mathbf{P} . Whatever is the complexity of φ , at one point in B , \mathbf{O} must have made an attack on some non-atomic assertion ψ made by \mathbf{P} such that ψ is a subformula of φ and ψ has only atoms (or an atom) as its immediate subformula(s). Otherwise, \mathbf{O} would still have a move to make and s would not be a winning strategy. However, we know that \mathbf{P} has answered this move by asserting some atomic formula p , and by Rule D10, this can be only in case if p has been previously asserted by \mathbf{O} . Thus, executing our algorithm will generate a tableau branch b' such that both $\mathbf{O}p$ and $\mathbf{P}p$ appear on it, which means that this branch b' is closed. The same holds of all other branches of s , so we know that the tableau corresponding to s must be closed. \square

Corollary 7. If φ is N-valid, then every N-tableau for $\mathbf{P}\varphi$ is closed.

Proof. Follows from Lemmas 12 and 14. \square

7.3. MAPPING TABLEAUX TO STRATEGIES.

We now turn to the direction from tableaux to strategies, and give a simple algorithm. We illustrate the algorithm with the example in Figure 1.

LEMMA 15. *If every N-tableau for $\mathbf{P}\varphi$ closes, then \mathbf{P} has a winning N-strategy for φ .*

Proof. Since there is no restriction on the order of application of the rules, it follows that if there is a closed tableau for $\mathbf{P}\varphi$, there is a closed, simple, \mathbf{O} -preferring tableau. We give an algorithm that will convert any closed, simple, \mathbf{O} -preferring tableau for $\mathbf{P}\varphi$ \mathcal{T} into a \mathbf{P} -winning N-strategy for φ .

1. Identify the closure atom in each branch; since the branches of the tableau close as soon as they can, we know the closure atom is unique.
2. Move the **P**-assertion of the closure atom so it is immediately after the **O**-assertion of the same atom. (See Figure 2.)




Figure 2. After step 2 has been applied.

3. Delete any **P**-assertions of non-closure atoms, if there are any.
4. Any time there are two **P**-nodes in a row, add a symbolic attack: The attack is ‘?’ after nodes of the form $\mathbf{P}(\varphi \vee \psi)$, $\wedge L$ at the beginning of any left tableau branch, and $\wedge R$ at the beginning of any right tableau branch. (See Figure 3.)




Figure 3. After step 4 has been applied.

5. Label all nodes with **P**-signed formulas as moves made by **P**; drop the **P** before these formulas.
6. Label all the remaining nodes as moves made by **O**; drop the **O** before the remaining signed formulas.
7. Number the nodes linearly in each branch, starting with 0; these numbers indicate the round where the move is made.

8. Assign stances and references as follows (see Figure 4):




Figure 4. After step 8 has been applied.

- a) The first node has stance “assert”, and no reference.
- b) If the node of the strategy was generated by an application of $\mathbf{P}\neg$, $\mathbf{O}\neg$, or $\mathbf{O}\rightarrow$ in the tableau, the stance is “attack” and the reference is n where n is the label of the node in the strategy corresponding to the premise of the rule application in the tableau.
- c) If the node of the strategy was generated by an application of $\mathbf{P}\rightarrow$ in the tableau, the stance is “defense” and the reference is n where n is the label of the node in the strategy corresponding to the application of $\mathbf{O}\rightarrow$ to the same premise in the tableau.
- d) If the node of the strategy was generated by an application of $\mathbf{P}\wedge$, $\mathbf{P}\vee_L$, or $\mathbf{P}\vee_R$ in the tableau, then in the strategy, the node is preceded by a symbolic attack; let n be the label of this symbolic attack. Label these nodes with the stance “defend” with reference n .
- e) If the node is a symbolic attack, the stance is “attack”, and the reference is the label of the previous node in the strategy.

Claim. The output of the algorithm is a winning N-strategy s for formula φ .

Proof. We must show (1) that s is an N-strategy; (2) that every branch ends with a \mathbf{P} -node, and (3) there are no legal \mathbf{O} moves by which any branches in s can be extended.

(1) The result is clearly an N-strategy; the second step assures that D10 is satisfied. D13 follows from the fact that any rule is only ever applied once, and hence \mathbf{O} does not attack a \mathbf{P} -formula more than once. The strategy begins with a \mathbf{P} -move, and the moves are alternating

because we always apply Rule $\mathbf{O} \rightarrow$ before Rule $\mathbf{P} \rightarrow$, and we have added symbolic attacks (moves by \mathbf{O}) between any two successive \mathbf{P} -nodes.

(2) Every branch ends with a \mathbf{P} -move, because in the tableau, either the branch closed when there was an \mathbf{O} -signed atom after the same atom \mathbf{P} -signed, and hence by step two of the algorithm, the branch of the strategy ends with a \mathbf{P} move; or the branch of the tableau closed when a \mathbf{P} -signed atom occurred after the same atom \mathbf{O} -signed had already occurred, and so this \mathbf{P} -assertion is the leaf.

(3) Towards a contradiction, suppose that one of s 's branches could be legally extended by an \mathbf{O} -move, either an attack or a defense.

1. \mathbf{O} 's move is a defense. There are two possibilities:

- a) It is a defensive move against a symbolic attack. This means that there must be a symbolic \mathbf{P} move somewhere in s . However, s was acquired from a tableau and the above algorithm does not introduce any symbolic attacks on \mathbf{P} 's part. Hence we know that there are no symbolic \mathbf{P} -moves in s . So we have a contradiction.
- b) It is a defensive move against an attack on implication. This means that at some node n in s \mathbf{O} has asserted a formula of the form $\psi \rightarrow \chi$ and at another node $m > n$ \mathbf{P} has attacked it by asserting ψ . Since s has been generated from a tableau by the above algorithm, we know that there are corresponding nodes n' and m' in the initial tableau such that at n' a formula of the form $\mathbf{O}(\psi \rightarrow \chi)$ is introduced and at node m' a formula of the form $\mathbf{P}\psi$ is introduced using $\mathbf{O}(\psi \rightarrow \chi)$ as a premise. However, there is no rule in Table III of this kind which means that s was not transformed from an N-tableau. Thus we have a contradiction.

2. \mathbf{O} 's move is an attack. There are five possibilities:

- a) It is an attack on a negation. This means there is a node n in s where \mathbf{P} has asserted a formula of the form $\neg\psi$. Since s was generated from an N-tableau by the above algorithm, we know that there is a corresponding node n' in the tableau whose label is $\mathbf{P}\neg\psi$. We know that there is no node m below n in s such that \mathbf{O} asserts ψ there, since by Rule D13 \mathbf{O} can only attack each move of \mathbf{P} once, and thus if he has already attacked \mathbf{P} 's assertion of $\neg\psi$, then it is not legal for him to attack it again, contra assumption. Since there is no such node in s , we know there is no node m' in \mathcal{T} that is labeled $\mathbf{O}\psi$, for

if there was, then by our algorithm, this would generate a node m in s labeled with \mathbf{O} 's assertion of ψ . But this contradicts our assumption that \mathcal{T} is simple, because if there is a node of the form $\mathbf{P}\neg\psi$ in \mathcal{T} , then the immediate next node will be $\mathbf{O}\psi$.

- b) It is an attack on an implication. This means that at some node n of s , \mathbf{P} has asserted a formula of the form $\psi \rightarrow \chi$ and it has not been attacked by \mathbf{O} . Given that s was transformed from an \mathbf{N} -tableau, we can be sure that there is corresponding node n' whose label is $\mathbf{P}(\psi \rightarrow \chi)$. We know that there is no node m below n in s such that \mathbf{O} asserts ψ there, since by Rule D13 \mathbf{O} can only attack each move of \mathbf{P} once, and thus if he has already attacked \mathbf{P} 's assertion of $\psi \rightarrow \chi$, then it is not legal for him to attack it again, contra assumption. Since there is no such node in s , we know there is no node m' in \mathcal{T} that is labeled $\mathbf{O}\psi$, for if there was, then by our algorithm, this would generate a node m in s labeled with \mathbf{O} 's assertion of ψ . But this contradicts our assumption that \mathcal{T} is simple, because if there is a node whose label has the form $\mathbf{P}(\psi \rightarrow \chi)$ in \mathcal{T} , then the immediate next node will be labeled by $\mathbf{O}\psi$.
- c) It is an attack on a disjunction, that is, the symbolic attack “?”. This means that there is some node n in s such that \mathbf{P} has asserted a disjunction $\psi \vee \chi$ at n and \mathbf{O} has not attacked it. However, our algorithm insures that each \mathbf{P} -assertion of a disjunction in s is followed by an attack on this disjunction on \mathbf{O} 's part. Thus we know that \mathbf{O} has attacked $\psi \vee \chi$ at node $n + 1$. Contradiction.
- d) It is an attack, asserting \wedge_L . This case is analogous to (c).
- e) It is an attack, asserting \wedge_R . This case is also analogous to (c).

Since in each case we have arrived at a contradiction, we can be sure that there is no move \mathbf{O} can make after the last \mathbf{P} -node. Thus we have proven (2), and we know that s is a winning \mathbf{N} -strategy. □

□

8. Conclusion

By making a simple and intuitive modification of the usual rules for classical dialogue games in the Felscher tradition, we obtained a set \mathbf{N} of dialogically valid formulas. We formulated the composition problem for a logic and related it to the question of what it means for a set of

formulas to be a logic. We resolved the composition problem for \mathbf{N} in the affirmative, allowing us to call \mathbf{N} a logic. Our positive solution to the composition problem was proved directly through semantic methods; that is, we worked solely with dialogue trees and strategies and did not need to follow the usual detour through cut-free sequent calculi or semantic tableaux. We have also provided a sound and complete tableaux system which gives a non-dialogical characterization of \mathbf{N} . We leave the problem of providing an independent axiomatization for \mathbf{N} , and questions of decidability and recursive axiomatizability, as future work.

The logic \mathbf{N} has a number of curious features, including a lack of unrestricted uniform substitution, and a failure to validate the implicational and conjunctive versions of *modus ponens* at the object-language level—despite the positive solution to its composition problem—which arise from the fact that if \mathbf{O} can defend once, he can always repeat this defense. As a result, \mathbf{N} privileges implications whose antecedents are atoms or negations, formulas which either can't be attacked or whose attacks cannot be defended against. The logic lies below \mathbf{CL} , but neither above nor below \mathbf{IL} , and is of interest because it is neither connexive nor relevant, i.e., it does not belong to either of the two families of well-known non-classical propositional logics which are not superintuitionistic.

It might be objected that if \mathbf{N} is a logic, it is a rather silly one.¹⁶ We do not dispute this; as it stands, there is no dialogically-independent reason to study \mathbf{N} , though this is not to say that some future application might not be found. Nevertheless, \mathbf{N} does illustrate how difficult the problems of providing justification and foundations for dialogical systems can be. If antecedent appeal to the resulting logic cannot be the primary motivation for combining a particular set of rules, then failing to obtain a serious logic cannot be an argument against the motivation for that combination. The set of rules generating \mathbf{N} is no less justified than the justification offered for $\mathbf{N} + \mathbf{E}$, which gives classical logic. Hence some kind of justification ought to be given for \mathbf{E} or for any set of rules containing \mathbf{E} , and this justification must be something other than “it gives us a nice logic”.

The results of our investigation are therefore significant in at least two respects. First, one of the important results in dialogical logic is the redundancy of the \mathbf{E} in the set of structural rules generating \mathbf{IL} . In this paper we investigated the role of \mathbf{E} in \mathbf{CL} , and showed that it is not similarly redundant. While we can safely drop \mathbf{E} from a ruleset known to capture intuitionistic logic, we cannot do so when working

with a ruleset CL known to capture classical logic, because we obtain N rather than CL.

This result should not be taken as a general statement about the necessity of E for classical dialogue semantics. Instead, it is better understood as analogous to results in the axiomatization of a certain class \mathcal{C} of structures where one finds that a certain axiom A must be included; this does not preclude the possibility that a slight adjustment of the class \mathcal{C} of structures could render the axiom A dispensable. Indeed, Felscher himself envisaged dialogues characterizing classical logic without the E rule (in [10, §3.10] he briefly considers so-called \mathcal{C} -dialogues in which the E rule is not available). Recent results show that E need not be assumed for the purpose of characterizing classical logic using dialogues [6, 40]. In [6], Clerbout uses a slightly different notion of dialogues (in effect, he works with a different class of structures) in which the notion of *rank*—numeric bounds on the number of repetitions the players are permitted—ensures that all dialogues whatsoever are finite, whereas Uckelman in [40] shows that, for Felscher dialogues, all that is necessary to obtain classical logic is to limit the number of repetitions that Opponent may make.

Second, our work highlights the importance of the composition problem in the dialogical setting, especially its relationship to the definition of ‘logic’ and the question of uniform substitution. Is it acceptable to drop unrestricted uniform substitution from a logic when there is no semantic (i.e., model-theoretic) motivations for doing so? Such questions when considered in the dialogical approach gain new relevance, showing the fruitfulness of the dialogical approach for philosophy of logic.

Acknowledgements

This paper is an extended version of [2], presented at the 4th Indian Conference on Logic and its Applications, Delhi, India, January 2011. The authors would like to thank the anonymous referees who provided many useful comments on earlier versions of this paper, as well as Shahid Rahman, whose suggestion prompted the results in §§6,7.

Notes

¹ For more on the history of dialogical logic, see [19, 22].

² Two of the current authors have discussed some of these matters further in [3]. The subject is also touched on in [29].

³ Note that it is not the case that every D-dialogue is an E-dialogue; an IL theorem that displays the divergence of the two is $\neg\neg(((\varphi \rightarrow \psi) \rightarrow \gamma) \rightarrow (((\psi \rightarrow \varphi) \rightarrow \gamma) \rightarrow \gamma))$.

⁴ Note that these two questions are not, strictly speaking, the same. In [40] a ruleset which generates CL but which does not have E is given.

⁵ See also the modal logics, mostly dynamic, in [15, fn. 1].

⁶ One can think of logics which do not satisfy unrestricted uniform substitution as being propositional theories, where the interpretation of the atomic variables is fixed in advance.

⁷ In many logics, it doesn't make sense to distinguish a schematic/substitutional notion of consequence from a non-substitutional one. However some logics, such as provability logics, do make such a distinction.

⁸ Thanks to one of the anonymous referees on a previous version of this paper for pointing this out to us.

⁹ This illustrates nicely Krabbe's point about the importance of structural rules regulating repetitive behavior ("the bugbear of dialogue theory") of the players [17, pp. 296, 303].

¹⁰ Another reason why **N** is not a relevance logic is because such formulas as $\neg(p \vee \neg p) \rightarrow q$ are valid. However, **N** could be considered a relevance logic in Rückert's sense of **P**-relevance logic, in which **P** must make all possible attacks and defenses. See [34, ch. 5].

¹¹ A similar thing happens in linear logic [8, §§2.1,6]; for a dialogical characterization of linear logic, see [4].

¹² More precisely, we attach to the end of $f(b)$ the result of cutting off the root of t , because the root of t and the leaf of b represent assertions by **P** of identical formulas.

¹³ Recall that the Gödel-Gentzen negative translation, however, does not preserve N-validity; see the proof of Lemma 11. This translation does not merely add negations, but also changes the shape of the formula.

¹⁴ Benedikt Löwe suggested this elegant solution.

¹⁵ We do not follow Felscher's developments of tableaux for intuitionistic logic, which are idiosyncractic for reasons related to his precise proof method.

¹⁶ Such an objection was raised by one of the anonymous referees.

References

1. Alama, J., A. Knoks, and S. L. Uckelman: 2011, 'Dialogue Games for Classical Logic (short paper)'. In: M. Giese and R. Kuznets (eds.): *TABLEAUX 2011: Workshops, Tutorials, and Short Papers*. pp. 82–86.
2. Alama, J. and S. L. Uckelman: 2010, 'A Curious Dialogical Logic and Its Composition Problem'. Preprint, <http://arxiv.org/abs/1008.0080>.
3. Alama, J. and S. L. Uckelman: 2012, 'What Is Dialogical About Dialogical Logic?'. In: H. J. Ribeiro (ed.): *Inside Arguments: Logic and the Study of Argumentation*. Cambridge Scholars Publication, pp. 207–222.
4. Blass, A.: 1992, 'A Game Semantics for Linear Logic'. *Annals of Pure and Applied Logic* **56**(1–3), 183–220.

5. Clerbout, N.: 2011, ‘Dialogical Games for First-Order Logic and Tableaux’. Presented at the 14e Congrès de Logique, Méthodologie et Philosophie des Sciences. 19–26 juillet 2011, Nancy, France.
6. Clerbout, N.: 2013, ‘First-Order Dialogical Games and Tableaux’. *Journal of Philosophical Logic*. Online, DOI 10.1007/s10992-013-9289-z.
7. D’Agostino, M.: 1999, ‘Tableau Methods for Classical Propositional Logic’. In: M. D’Agostino, D. M. Gabbay, R. Hähnle, and J. Posegga (eds.): *Handbook of Tableau Methods*. Springer, pp. 45–123.
8. Di Cosmo, R. and D. Miller: 2010, ‘Linear Logic’. In: E. N. Zalta (ed.): *Stanford Encyclopedia of Philosophy*. CSLI Publications, fall edition. <http://plato.stanford.edu/archives/fall2010/entries/logic-linear/>.
9. Felscher, W.: 1985, ‘Dialogues, Strategies, and Intuitionistic Provability’. *Annals of Pure and Applied Logic* **28**, 217–254.
10. Felscher, W.: 1986, ‘Dialogues as Foundation for Intuitionistic Logic’. In: D. Gabbay and F. Günther (eds.): *Handbook of Philosophical Logic*, Vol. III. Reidel, pp. 341–372.
11. Fermüller, C. and A. Ciabatonni: 2003, ‘From Intuitionistic Logic to Gödel-Dummett Logic via Parallel Dialogue Games’. In: *ISMVL 2003, Proceedings of the 33rd IEEE International Symposium on Multiple-Valued Logic, May 16–19 2003, Tokyo, Japan*. pp. 188–195.
12. Fermüller, C. G.: 2003, ‘Parallel Dialogue Games and Hypersequents for Intermediate Logics’. In: M. C. Mayer and F. Pirri (eds.): *TABLEAUX 2003 Automated Reasoning with Analytic Tableaux and Related Methods*. pp. 48–64.
13. Finger, M. and D. Gabbay: 2006, ‘Cut and Pay’. *Journal of Logic, Language and Information* **15**(3), 195–218.
14. Hodges, W.: 2009, ‘Logic and Games’. In: E. N. Zalta (ed.): *Stanford Encyclopedia of Philosophy*. CSLI Publications, Spring edition. <http://plato.stanford.edu/archives/spr2009/entries/logic-games/>.
15. Holliday, W. H., T. Hoshi, and T. F. Icard III: 2012, ‘A Uniform Logic of Information Dynamics’. In: T. Bolander, T. Braüner, S. Ghilardi, and L. Moss (eds.): *Advances in Modal Logic*, Vol. 9. pp. 348–367.
16. Keiff, L.: 2009, ‘Dialogical Logic’. In: E. N. Zalta (ed.): *Stanford Encyclopedia of Philosophy*. CSLI Publications, Summer edition. <http://plato.stanford.edu/archives/sum2009/entries/logic-dialogical/>.
17. Krabbe, E. C. W.: 1985, ‘Formal Systems of Dialogue Rules’. *Synthese* **63**, 295–328.
18. Krabbe, E. C. W.: 2001, ‘Dialogue Foundations: Dialogue Logic Revisited [sic]’. *The Aristotelian Society Supplement* **75**, 33–49.
19. Krabbe, E. C. W.: 2006, ‘Dialogue Logic’. In: D. M. Gabbay and J. Woods (eds.): *Handbook of the History of Logic*, Vol. 7. Elsevier, pp. 665–704.
20. Lorenz, K.: 1961, ‘Arithmetik und Logik als Spiele’. Ph.D. thesis, Universität Kiel.
21. Lorenz, K.: 1968, ‘Dialogspiele als semantische Grundlage von Logikkalkülen’. *Archiv Math. Logik Grundlagenforsch.* **11**, 32–55, 73–100.
22. Lorenz, K.: 2001, ‘Basic Objectives of Dialogue Logic in Historical Perspective’. *Synthese* **127**, 255–263.
23. Lorenzen, P.: 1955, *Einführung in die operative Logik und Mathematik*. Springer.
24. Lorenzen, P. and K. Lorenz: 1978, *Dialogische Logik*. Wissenschaftliche Buchgesellschaft.

25. Marcos, J.: 2009, 'What is a Non-truth-functional Logic?'. *Studia Logica* **92**, 215–240.
26. McCall, S.: 1996, 'Connexive Implication'. *Journal of Symbolic Logic* **31**(3), 415–433.
27. Priest, G.: 1999, 'Negation as Cancellation, and Connexive Logic'. *Topoi* **18**, 141–148.
28. Priest, G.: 2006, *In Contradiction: A Study of the Transconsistent*. Oxford University Press, 2nd edition.
29. Rahman, S.: 2011, 'Negation in the Logic of First Degree Entailment and Tonk: A Dialogical Study'. In: G. Primeiro and M. Marion (eds.): *The Realism-Antirealism Debate in the Age of Alternative Logics*. Springer, pp. 175–202.
30. Rahman, S., N. Clerbout, and L. Keiff: 2009, 'On Dialogues and Natural Deduction'. In: G. Primeiro and S. Rahman (eds.): *Acts of Knowledge: History, Philosophy and Logic*. London: College Publications, pp. 301–355.
31. Rahman, S. and L. Keiff: 2004, 'On How to be a Dialogician'. In: D. Vanderveken (ed.): *Logic, Thought and Action*. Springer, pp. 359–408.
32. Rescher, N. and R. Manor: 1970, 'On Inference From Inconsistent Premises'. *Theory and Decision* **1**(2), 179–217.
33. Restall, G.: 2009, 'Substructural Logics'. In: E. N. Zalta (ed.): *Stanford Encyclopedia of Philosophy*. CSLI Publications, summer edition. <http://plato.stanford.edu/archives/sum2009/entries/logic-substructural/>.
34. Rückert, H.: 2007, 'Dialogues as a Dynamic Framework for Logic'. Ph.D. thesis, Universiteit Leiden.
35. Smullyan, R. M.: 1995, *First-Order Logic*. Dover.
36. Sørensen, M. H. and P. Urzyczyn: 2006, *Lectures on the Curry-Howard Isomorphism*. Eslevier.
37. Sørensen, M. H. and P. Urzyczyn: 2007, 'Sequent calculus, Dialogues, and Cut Elimination'. In: E. Barendsen, H. Geuvers, V. Capretta, and M. Niqui (eds.): *Reflectons on Type Theory, λ -Calculus, and the Mind: Essays Dedicated to Henk Barendregt on the Occasion of His 60th Birthday*. Radboud Universiteit Nijmegen, pp. 253–261.
38. Tarski, A.: 1983, 'On the Concept of Logical Consequence'. In: J. Corcoran (ed.): *Logic, Semantics, Metamathematics*. Hackett, 2 edition, pp. 409–420. Translated by J. H. Woodger.
39. Troelstra, A. S. and H. Schwichtenberg: 2000, *Basic Proof Theory*, Vol. 43 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2nd edition.
40. Uckelman, S. L.: 2012, 'Some Remarks on the E-Rule in Dialogical Logic'. Preprint, in submission.